

DOCUMENT RESUME

ED 039 102

24

RE 002 719

AUTHOR Calfee, Robert C.
TITLE The Role of Mathematical Models in Optimizing Instruction.
INSTITUTION Wisconsin Univ., Madison. Research and Development Center for Cognitive Learning.
SPONS AGENCY Office of Education (DHEW), Washington, D.C. Bureau of Research.
REPORT NO TP-17
BUREAU NO RR-5-0216
PUB DATE Mar 69
CONTRACT OEC-5-10-154
NOTE 32p.

EDRS PRICE MF-\$0.25 HC-\$1.35
DESCRIPTORS Computer Assisted Instruction, *Decision Making, *Instructional Design, *Instructional Improvement, *Learning Processes, *Mathematical Models, Verbal Learning

ABSTRACT

Consideration of computer-assisted instruction in the classroom has led to an analysis of the educational process including the need for developing more adequate models of the learning processes and increased attention to the function of the teacher (human or computer) as a decision maker. Given a descriptive model of the learning process, it is in some cases possible to derive optimal conditions for instruction. The technique of optimization is illustrated by considering (1) the optimal number of blocks into which a list of items should be divided (whole versus part) and (2) optimal choice of items as a function of the previous response history, which involves dropping "learned" items during a training sequence. The efficiency of the dropout procedure depends on list criterion only. A learning model which serves as a plausible approximation was developed for both the above examples. The problem of relative efficiency is considered. Tables, appendixes, and references are included. (WB)

THE ROLE OF MATHEMATICAL MODELS IN OPTIMIZING INSTRUCTION

U.S. DEPARTMENT OF HEALTH, EDUCATION
& WELFARE

OFFICE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED
EXACTLY AS RECEIVED FROM THE PERSON OR
ORGANIZATION ORIGINATING IT. POINTS OF
VIEW OR OPINIONS STATED DO NOT NECES-
SARILY REPRESENT OFFICIAL OFFICE OF EDU-
CATION POSITION OR POLICY

WISCONSIN RESEARCH AND DEVELOPMENT

**CENTER FOR
COGNITIVE LEARNING**



ED039102

Theoretical Paper No. 17

THE ROLE OF MATHEMATICAL MODELS IN
OPTIMIZING INSTRUCTION

By Robert C. Calfee

Report from the Project on Language Concepts and Cognitive Skills
Related to the Acquisition of Literacy
Robert C. Calfee and Richard L. Venezky, Principal Investigators

Wisconsin Research and Development
Center for Cognitive Learning
The University of Wisconsin
Madison, Wisconsin

March 1969

This research was performed pursuant to Grant MH 12637 from the National Institute of Mental Health and to a contract with the United States Office of Education, Department of Health, Education, and Welfare, under the provisions of the Cooperative Research Program at The University of Wisconsin Research and Development Center.

Center No. C-03, Contract OE 5-10-154

NATIONAL EVALUATION COMMITTEE

Samuel Brownell
Professor of Urban Education
Graduate School
Yale University

Henry Chauncey
President
Educational Testing Service

Elizabeth Koontz
President
National Education Association

Patrick Suppes
Professor
Department of Mathematics
Stanford University

Launor F. Carter
Senior Vice President on
Technology and Development
System Development Corporation

Martin Deutsch
Director, Institute for
Developmental Studies
New York Medical College

Roderick McPhee
President
Punahou School, Honolulu

***Benton J. Underwood**
Professor
Department of Psychology
Northwestern University

Francis S. Chase
Professor
Department of Education
University of Chicago

Jack Edling
Director, Teaching Research
Division
Oregon State System of Higher
Education

G. Wesley Sowards
Director, Elementary Education
Florida State University

UNIVERSITY POLICY REVIEW BOARD

Leonard Berkowitz
Chairman
Department of Psychology

John Guy Fowlkes
Director
Wisconsin Improvement Program

Herbert J. Klausmeier
Director, R & D Center
Professor of Educational
Psychology

M. Crawford Young
Associate Dean
The Graduate School

Archie A. Buchmiller
Deputy State Superintendent
Department of Public Instruction

Robert E. Grinder
Chairman
Department of Educational
Psychology

Donald J. McCarty
Dean
School of Education

***James W. Cleary**
Vice Chancellor for Academic
Affairs

H. Clifton Hutchins
Chairman
Department of Curriculum and
Instruction

Ira Sharkansky
Associate Professor of Political
Science

Leon D. Epstein
Dean
College of Letters and Science

Clauston Jenkins
Assistant Director
Coordinating Committee for
Higher Education

Henry C. Weinlick
Executive Secretary
Wisconsin Education Association

EXECUTIVE COMMITTEE

Edgar F. Borgatta
Birmingham Professor of
Sociology

Russell J. Hosler
Professor of Curriculum and
Instruction and of Business

Wayne Otto
Professor of Curriculum and
Instruction (Reading)

Richard L. Venezky
Assistant Professor of English
and of Computer Sciences

Max R. Goodson
Professor of Educational Policy
Studies

***Herbert J. Klausmeier**
Director, R & D Center
Professor of Educational
Psychology

Robert G. Petzold
Associate Dean of the School
of Education
Professor of Curriculum and
Instruction and of Music

FACULTY OF PRINCIPAL INVESTIGATORS

Ronald R. Allen
Associate Professor of Speech
and of Curriculum and
Instruction

Gary A. Davis
Associate Professor of
Educational Psychology

Max R. Goodson
Professor of Educational Policy
Studies

Richard G. Morrow
Assistant Professor of
Educational Administration

Vernon L. Allen
Associate Professor of Psychology
(On leave 1968-69)

M. Vere DeVault
Professor of Curriculum and
Instruction (Mathematics)

Warren O. Hagstrom
Professor of Sociology

Wayne Otto
Professor of Curriculum and
Instruction (Reading)

Nathan S. Blount
Associate Professor of English
and of Curriculum and
Instruction

Frank H. Farley
Assistant Professor of
Educational Psychology

John G. Harvey
Associate Professor of
Mathematics and Curriculum
and Instruction

Milton O. Pella
Professor of Curriculum and
Instruction (Science)

Robert C. Calfee
Associate Professor of Psychology

John Guy Fowlkes (Advisor)
Professor of Educational
Administration
Director of the Wisconsin
Improvement Program

Herbert J. Klausmeier
Director, R & D Center
Professor of Educational
Psychology

Thomas A. Romberg
Assistant Professor of
Mathematics and of
Curriculum and Instruction

Robert E. Davidson
Assistant Professor of
Educational Psychology

Lester S. Golub
Lecturer in Curriculum and
Instruction and in English

Burton W. Kreitlow
Professor of Educational Policy
Studies and of Agricultural
and Extension Education

Richard L. Venezky
Assistant Professor of English
and of Computer Sciences

MANAGEMENT COUNCIL

***Herbert J. Klausmeier**
Director, R & D Center
Acting Director, Program 1

Thomas A. Romberg
Director
Programs 2 and 3

James E. Walter
Director
Dissemination Section

Dan G. Woolpert
Director
Operations and Business

Mary R. Quilling
Director
Technical Section

* COMMITTEE CHAIRMAN

STATEMENT OF FOCUS

The Wisconsin Research and Development Center for Cognitive Learning focuses on contributing to a better understanding of cognitive learning by children and youth and to the improvement of related educational practices. The strategy for research and development is comprehensive. It includes basic research to generate new knowledge about the conditions and processes of learning and about the processes of instruction, and the subsequent development of research-based instructional materials, many of which are designed for use by teachers and others for use by students. These materials are tested and refined in school settings. Throughout these operations behavioral scientists, curriculum experts, academic scholars, and school people interact, insuring that the results of Center activities are based soundly on knowledge of subject matter and cognitive learning and that they are applied to the improvement of educational practice.

This Theoretical Paper is from the Language Concepts and Cognitive Skills Related to the Acquisition of Literacy Project in Program 1. General objectives of the Program are to generate new knowledge about concept learning and cognitive skills, to synthesize existing knowledge, and to develop educational materials suggested by the prior activities. Contributing to these Program objectives, this project's basic goal is to determine the processes by which children aged four to seven learn to read and to identify the specific reasons why many children fail to acquire this ability. Later studies will be conducted to find experimental techniques and tests for optimizing the acquisition of skills needed for learning to read.

CONTENTS

	Page
Abstract	vii
I. Introduction	1
II. Static Decision Strategies	4
III. Dynamic Decision Strategies	7
IV. Efficiency of Decision Strategies	11
V. Heterogeneity of Learning Rate Parameters	14
VI. Conclusions and Implications	17
Appendix A: Evaluation of Block-Size Function	19
Appendix B: Derivation of <i>a posteriori</i> Probability of State <i>L</i> in All-or-None Model	21
Appendix C: Conditions for Optimality of Largest-Immediate- Gain Strategy	22
References	27

LIST OF TABLES AND FIGURE

Table	Page
1 Proportion of Items in <i>L</i> at End of Training Following Standard, Dropout, and Error-Proportion Decision Strategies	11
2 Mean State Values at End of 100 Presentations, 10-Item List for Various Decision Strategies, All-or-None and Incremental Models with Heterogeneity of Parameters	16

Figure	
1 Flow Diagram of Instructional Sequence	3

ABSTRACT

With the advent of computers in the classroom, the decision structure implicit in the teaching process has been made obvious. Given a descriptive model of the learning process, it is in some cases possible to derive optimal conditions for instruction.

The technique of optimization is illustrated by two examples. In the first, the question of the optimal block-size for the learning of foreign language vocabulary is discussed. The prediction from a two-operator learning model is that presentation of the entire list should produce the most efficient learning. Despite certain failings of the model, available data are consistent with the prediction.

The second example concerns the application of *dynamic programming* to optimal selection of items in a paired-associate problem. It is shown that, given some general constraints on the learning model and the gain function relating the increase in probability of a correct response to the model, the optimal decision at any point is to present that item from the list which produces the largest immediate increment in response probability.

The problem of relative efficiency is also taken up, and it is shown that, assuming that learning occurs in an all-or-none fashion, a "dropout" procedure is considerably more efficient than the standard procedure commonly used if presentation of items stops only after a list criterion is reached. The efficiency of the dropout procedure is shown to depend only on the list criterion.

I INTRODUCTION

The digital computer has begun to take its place in the classroom as an instructional device. Presently, such installations are limited mainly to experimental settings (Suppes, 1966). Although it remains to be seen whether computer-aided instructional systems will prove either pedagogically or economically practical, the problem of efficient utilization of such systems is likely to hold the interest of psychologists and other behavioral scientists for some time.

The computer may be used simply as another means of implementing a traditional instructional format; e.g., one might simply load a programmed text into the machine, which could then present the material to the student. A more exciting potential of computer-aided instruction concerns the individualization of instruction, a process in which the machine determines each step in the instructional sequence on the basis of past performance using a suitable decision strategy.

A couple of examples might be helpful. One decision strategy might entail selection of a presentation rate which best matches the student's ability to assimilate information. Suppose some material is presented and the student is then tested. If he responds correctly, the computer goes on to new material; if incorrectly, the problem is presented again, perhaps in a different form. Thus new material is presented only as fast as the student can assimilate it.

An elaboration of this notion involves deciding what "track" to put a student on. In the teaching of beginning reading in American elementary schools, a three-track system is typically used. Classes are divided into children of high, middle, or low reading ability. Low ability readers use basal primers with restricted vocabularies where the stress is on rote memorization of words, while for high ability readers the vocabularies tend to be much larger and there is more emphasis on development of word attack skills

based on an analysis of letter-sound relations. Computer-aided systems could in principle continuously evaluate a child's performance and change him from track to track as appropriate. If a child were having a great deal of trouble during the early stages of reading, he would be placed on a low track. As he began to perform more adequately, the machine could take the necessary steps to move him to the next higher track.

It is easy to imagine in a general way how computer-aided systems might be used to institute significant innovations in education. On the other hand, the consideration of such systems has led to analyses of the educational process which reveal some fundamental questions. It seems safe to say that if a revolution in education is brought about in this generation, it will result not from the introduction of computers into the classroom, but from (a) *a more detailed analysis of specific curriculum goals*, (b) *development of more adequate models of the process by which students learn*, and (c) *increased attention to the function of the teacher (human or computer) as a decision maker*. The first of these matters is a problem for curriculum specialists; the second and third problems are the province of the psychologist. Each of these matters constitutes a fundamental and unsolved problem in education.

For example, the model of the student implicit in most educational practices may perhaps be best stated as "practice makes perfect." The decision rule has been, "if at first you don't succeed, try again." The failing student is encouraged to repeat the same activity until either he "gets it" or runs out of time. The adequacy of both the model and the decision rule may be called into question.

When one considers the variety of combinations of ideas about what is to be taught, how it is going to be learned, and what decision policies will be most efficient, it is obvious that the odds are very small that optimal teaching programs will be arrived at

by chance. Yet the absence of formal procedures for the development of teaching programs and the strong reliance on intuition, common sense, and previous practice suggest that just such a search for a needle in a haystack has been going on.

Some efforts at formalization have appeared during the past few years within a specific theoretical framework. Suppose that for a particular portion of a curricular program (e.g., learning the names of the letters of the alphabet), there exists an adequate mathematical model of the learning process. Then in some instances, this model may allow one to discover unique decision rules which are optimal in the sense that the student will learn as much as possible in a fixed period of time.

In this paper, two optimization problems will be considered: (a) the optimal number of blocks into which a list of items should be divided (the whole versus part problem) and (b) optimal choice of items as a function of the previous response history, which involves the question of dropping out "learned" items during a training sequence. For both these examples, it has been possible to obtain solutions to the optimality problem by means of a learning model which serves as a plausible approximation, albeit a gross simplification. In neither example are there sufficient data for empirical evaluation of the efficiency of various presentation strategies. Nevertheless, these two problems will illustrate the utility of mathematical learning models in the derivation of optimal teaching procedures, and the techniques as well as the problems of this approach.

Before proceeding, it may be useful to characterize the instructional paradigm more explicitly. First, the material to be taught will be conceptualized as a set of stimulus items, S_i , for each of which there exist one or more correct responses, R_i . These S-R pairs may be as simple as saying the phoneme /e/ to the symbol A, or as complex as writing an answer to the question, "What were the major causes of the American Civil War?" after reading a 1000-word passage on the topic. We also define a history at any point in time as the knowledge available to the instructor at that time about the preceding sequence of S and R events. Then, following Groen and Atkinson (1966), the succession of events in an instructional system can be diagrammed as in Fig. 1. Each presentation begins with the selection of some S_i , followed by some response, then the selection of the next S_j , and so forth. It may be easiest to think of this procedure in terms of the learning of a foreign language vocabulary, such as acquiring the English equivalents of German words. Generally speaking, it will be assumed that there are unique responses to be associated with each stimulus, and that each pair must be presented several times for learning to take place. If the selection is contingent on the previous response history, the system will be referred to as *dynamic* (response-sensitive in Groen and Atkinson's terms) while if the selection is independent of the history the system will be called *static* (or response-insensitive).

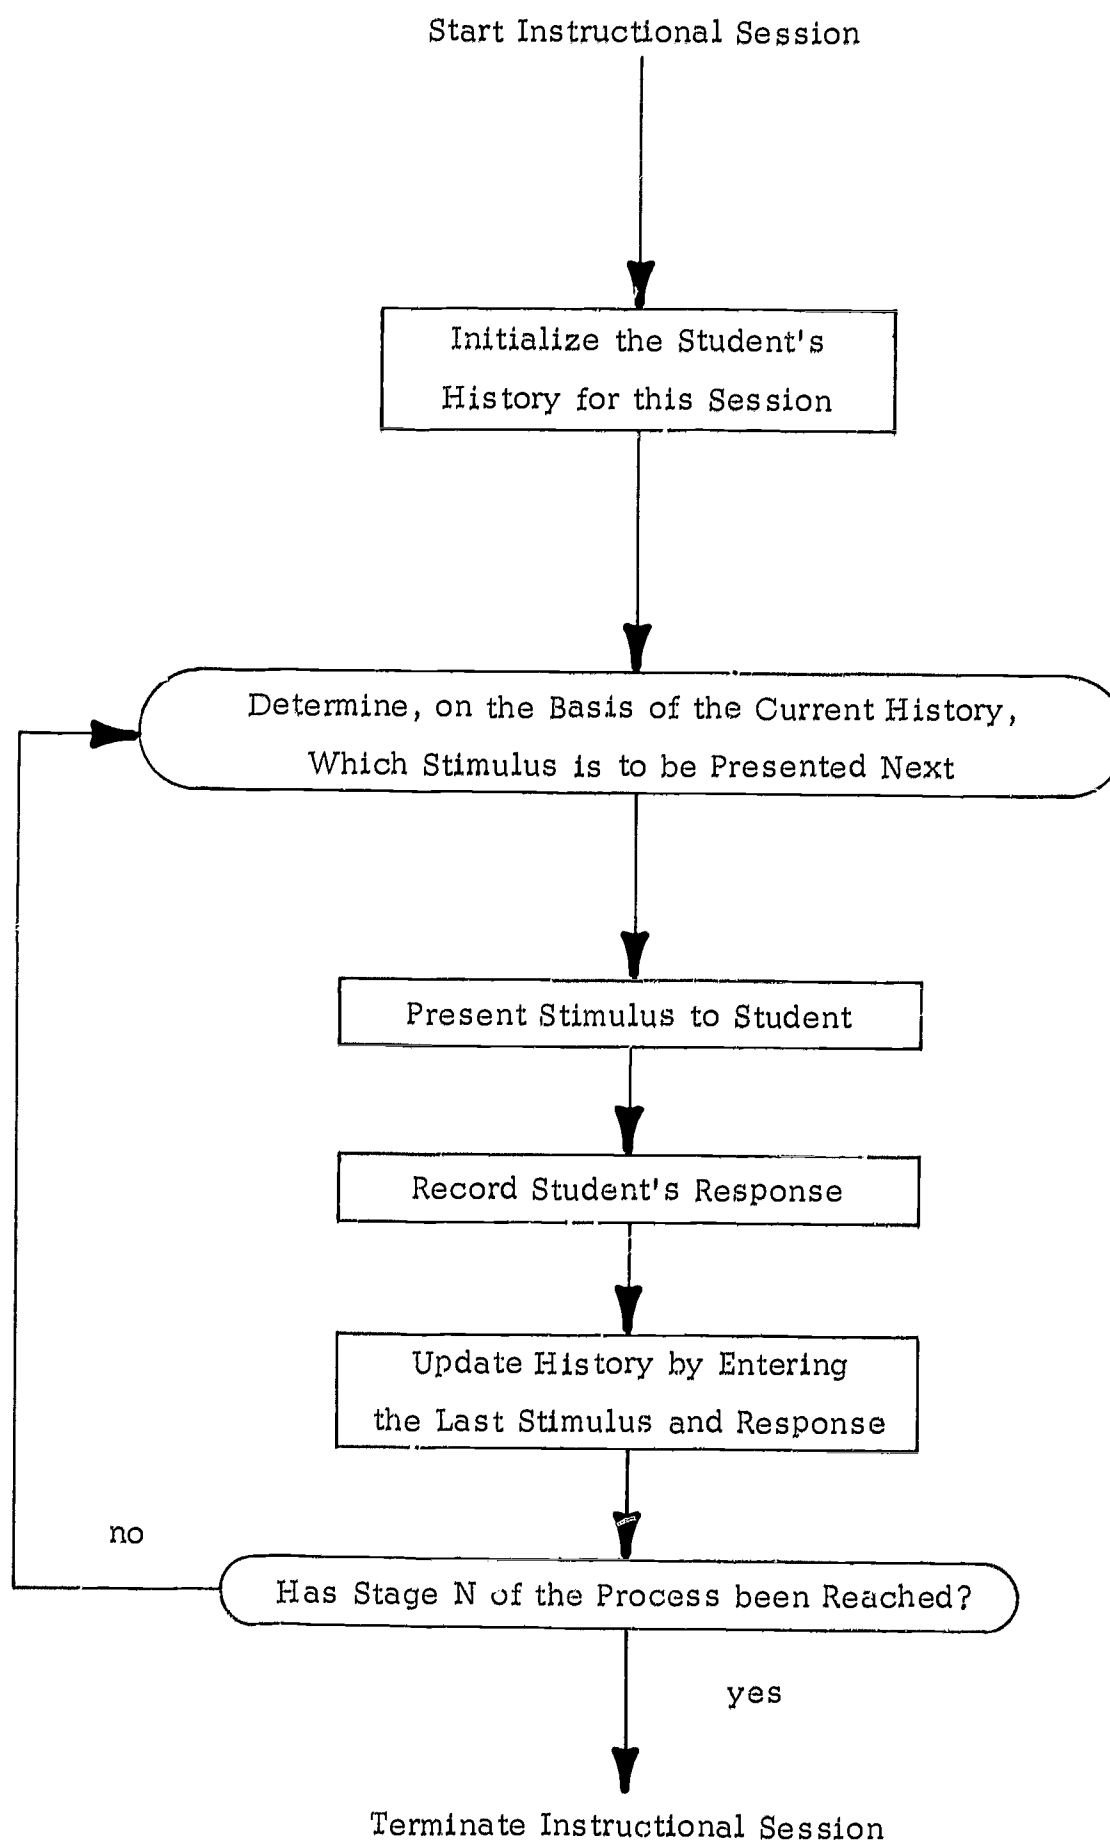


Figure 1. Flow Diagram of Instructional Sequence

II STATIC DECISION STRATEGIES

The major feature of static strategies is that the sequence of S-R items can be prepared prior to the beginning of instruction; the history is not necessary. As an example of a static problem, suppose the number of items in a list is quite large so that it appears sensible to break it into two or more sublists or blocks, each of which is taught in turn. What is the optimum number of blocks or, equivalently, the optimum block size?

We will consider Suppes' (1964) analysis of this problem in some detail, because a solution exists and is relatively easy to present. The problem can be specified as follows. A list of M items is to be divided into blocks of k items each. Each block is given N trials (i.e., the items in the first block are each presented N times, then the items in the second block, etc.), and then a final test is administered covering the entire list. The problem is to find that value of k which maximizes the proportion of correct responses on the final test.

Suppes assumed that in this situation the learning process could be described by the following *two-operator linear model*.¹ Consider the simplest possible case. Two items are to be learned; i.e., $M = 2$. Each time item 1 is presented and studied by the student, the probability of error on the next test of that item is assumed to decrease, whereas each time item 2 is presented for study, the probability of an incorrect response for item 1 increases. In short, study presentations for an item lead to learning of it, and presentations of other items lead to forgetting of it. The decrements and increments in error probability are assumed to be proportional to the

amount remaining to be learned or forgotten, respectively. Specifically, if $q_{i,n+1}$ is the error probability for item i on the $n+1$ st presentation, then the effect of studying i on n is

$$q_{i,n+1} = a q_{i,n}, \quad 0 < a < 1. \quad (1)$$

Suppose $p_{i,n+1}$ is the success probability for i on $n+1$. Then the effect of studying some other item on n is

$$p_{i,n+1} = b p_{i,n}, \quad 0 < b < 1. \quad (2)$$

Since by definition $p_{i,n+1} = 1 - q_{i,n+1}$ by simple algebra it can be seen that (2) may be written as $q_{i,n+1} = (1 - b) + b q_{i,n}$ which will prove to be a more useful form. The parameters a and b constitute the learning and forgetting rates which depend on the student's ability, the material to be learned, and other factors. As it turns out, the solution to the optimization problem depends only on which is larger, a or b .

Equations 1 and 2 (the two operators) are linear difference equations, the properties of which are well known (Goldberg, 1961). In particular, for an equation of the general form

$$x_{t+1} = R x_t + S, \quad (3)$$

the solution is

$$x_{t+1} = R^t x_1 + S \left[\frac{1 - R^t}{1 - R} \right]. \quad (4)$$

As an example, suppose item i is studied for t consecutive trials. For convenience, assume that initially the chance of giving the correct answer is zero, i.e., that $q_{i,1} = 1$.

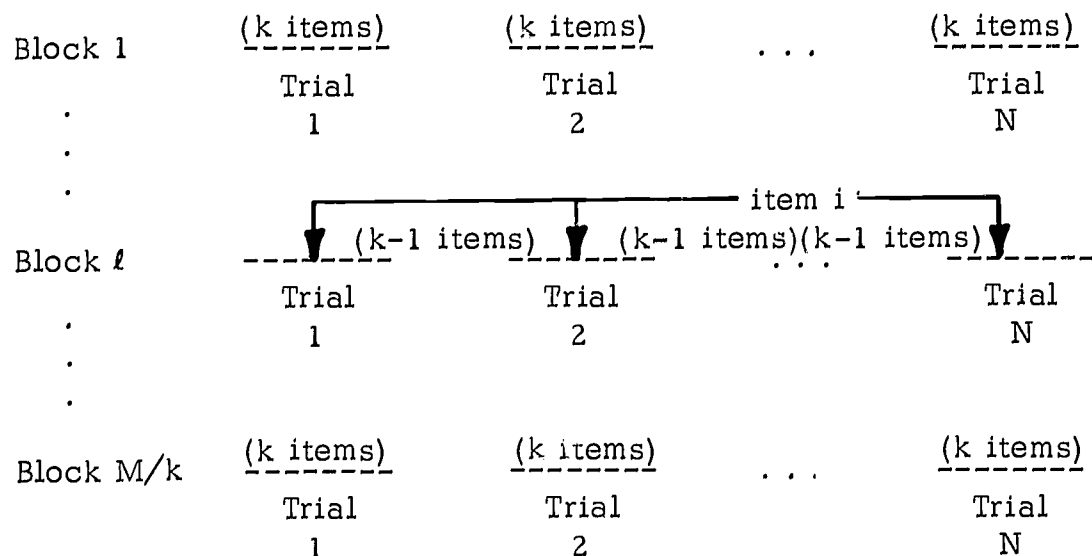
Applying (1), $q_{i,2} = a q_{i,1} = a$; $q_{i,3} = a q_{i,2} = a^2$; \dots ; $q_{i,t+1} = a^t$, which agrees with the solution given in (4) for $a = R$ and $S = 0$.

We will now present in outline form Suppes' solution to the problem of optimal

¹In the sections that follow, where derivations involving specific models of learning are carried out, it is assumed that the reader has some background in probability theory of familiarity with the literature at the level of Atkinson, Bower, and Crothers (1965).

block size based on this model. For convenience, it may be assumed that $q_{i,1} = 1$ for all i . Suppose the list is divided into blocks of

size k . Taking a particular item i , at some point the block containing i receives its series of N trials, as indicated in the diagram below:



The analysis is simplified by assuming that an item is presented at the same relative location in the block on each of the N trials. There will be $N-1$ complete cycles on each of which i is presented followed by the $k-1$ other items in the block. Then i receives its N th presentation, the remaining items in the block are presented, and then there are Nk presentations for each of the blocks remaining to be studied. Let R_i denote the number of presentations of other items following the last presentation of i .

Consider the events of a single cycle. Let $q_i^{(j)}$ be the probability of an error for i at the time of its presentation on trial j . After i is presented and studied, the new probability will be $aq_i^{(j)}$. Following the $k-1$ other items in the block (i.e., just prior to its presentation on trial $j+1$), from (2) and (4)

$$q_i^{(j+1)} = b^{k-1} (aq_i^{(j)}) + \left[\frac{(1-b)(1-b^{k-1})}{(1-b)} \right] = ab^{k-1} q_i^{(j)} + (1-b^{k-1}). \quad (5)$$

However, (5), which incorporated all the events in a cycle, is itself a linear difference equation. Hence, following the $N-1$ cycles in a block, or just prior to the last presentation

$$q_i^{(N)} = (ay)^{N-1} + \frac{(1-y)[1-(ay)^{N-1}]}{1-ay} \quad (6)$$

where $y = b^{k-1}$. There follows the N th presentation of i , and then R_i presentations of other items, whence the probability of an error at the time of the final test is

$$q_i^{(F)} = ab^{R_i} q_i^{(N)} + (1-b^{R_i}). \quad (7)$$

Only the quantity R_i varies from item to item.

The calculation of the average of $q_i^{(F)}$ over items is not trivial, but the details are not particularly instructive (but cf. Suppes, 1964). Letting $v = a/b$ and $w = b^k$, it can be shown that the expectation of $q_i^{(F)}$ is

$$E(q) = 1 - \left[\frac{(1-b^{MN})(1-a)}{(1-b)M} \right] \left[\frac{1-(vw)^N}{1-vw} \right] \left[\frac{1-w}{1-w^N} \right]. \quad (8)$$

Only the last two terms in (8) need to be considered in determining the optimal choice of k , since only w is a function of k . Notice that if $v = 1$ (i.e., if $a = b$) then the last two terms cancel. Hence, if the rate parameters are equal (if the relative amount learned during study of an item is exactly offset by presentation of some other item), performance does not depend on block size, so there is no optimal block size.

If $v < 1$, then $E(q)$ decreases monotonically with k (cf. Appendix A) and so the largest block size should be used— $k_{opt} = M$. Note that $v < 1$ implies that $a < b$, so that the rate of learning when an item is presented for

study is greater than the rate of forgetting when some other item is presented. Conversely, if $v > 1$, $E(q)$ increases with k , hence $k_{\text{opt}} = 1$.

According to this model, there are only two optimal block sizes—either take the entire list and present it as a single block N times, or present the first item N times, then the second item, etc.—and the choice depends only on whether a or b is greater. In fact, only the case where $a < b$ is of more than passing interest. It is not difficult to show that if $b < a$ then letting $k = 1$ (which is optimal)

$$\lim_{N \rightarrow \infty} E(q) = 1 - \frac{1}{M} \quad (10)$$

In other words, only one item will be learned. Since in most learning situations perfect performance is reached after prolonged practice, it appears reasonable to assume $a > b$. One would therefore predict that *whole* learning, as it has been referred to in the psychological literature, should be more efficient (or at least as efficient) than *part* learning, in which the list is broken into some number of blocks. However, the relative efficiency of a part procedure will depend on the list length and the number of blocks into which the list is broken, among other factors.

The experimental evidence on the point is not especially enlightening. Reports can be cited (cf. Osgood, 1953, pp. 540 ff.) in which the whole procedure was superior, or the part procedure was superior, or (more usually) no difference was obtained. Unpublished work in our laboratory involving relatively long lists of German-English pairs, and the extensive investigations by Crothers and Suppes (1967) on acquisition of Russian vocabulary indicate that a list may be broken into two or three blocks with no loss in efficiency, and in some instances a gain in efficiency is observed (Exps. X and XI, Crothers & Suppes, 1967). Breaking a list into a large number of small blocks generally seems detrimental, so that under certain conditions the empirically optimal block size may be an intermediate value between 1 and M , contrary to the predictions of the two-operator model.

This model is plainly an oversimplification of those processes which allow a student to learn a paired-associate list. Certain features of the model compromise its usefulness for prediction of optimal teaching conditions in all but a few contexts. For example, in (10) it was shown that if $a < b$,

at most one item would be recalled after a large number of trials. This prediction arises because no matter how many study trials are given to an item, it is *completely* forgotten after the presentation of a large number of other items. Consider an alternative model where, in addition to the short-term acquisition and retention processes described by (1) and (2), more permanent storage of information takes place during study trials. This permanent storage can be expressed by the following modification of (2):

$$q_{i,n+1} = bq_{i,n} + (1-b)\gamma^r \quad (11)$$

The parameter γ is the rate of permanent storage, and r is the number of preceding study trials on i . Suppose $q_i^{(r)}$ is the error probability for i following the r study trials. As k , the number of other items following the r^{th} presentation of i , becomes large, from (1) and (11)

$$q_i^{(r)} = \lim_{k \rightarrow \infty} b^k (aq_i^{(r-1)}) + (1-b^k)\gamma^{r-1} = \gamma^{r-1}. \quad (12)$$

The parameter a now effectively represents the probability that an incorrect response will occur immediately following a study trial. Considerable evidence suggests that this probability is small. Making the assumption that a is zero, it is not difficult to show that

$$E(q) = \gamma^N \left\{ 1 - \left[\frac{(1-b)^{MN}}{(1-b)M} \right] \left[\frac{1-w}{1-w^N} \right] \right\} \quad (13)$$

The derivative of (13) with regard to k is always negative, implying that the whole-list procedure will always prove most efficient.

Another shortcoming in both models is that they assume only the simplest interactions among items within a list. However, there are a number of plausible sources of interaction: the similarity between the elements of i and j , the degree of learning of an item, and the extent to which the student decides to attend carefully to one particular item rather than another.

Nonetheless, this example illustrates the technique of optimization, and the kinds of applications which are possible. Extension of these techniques to matters of the sort just mentioned await the development of more suitable and comprehensive models.

III DYNAMIC DECISION STRATEGIES

In this section we consider optimization problems which conform to the same S-R paradigm as before, but where the student's history in the session is used in choosing the sequence of presentation. The earliest paper on this topic is by Smallwood (1962). He posed the optimization question in a general form and introduced the technique of dynamic programming into the learning literature. His ideas about learning models were not as well developed, and his paper will not be considered here in any detail.

Instead, we will look at a question considered recently by several investigators (Groen & Atkinson, 1966; Karush & Dear, 1966; Matheson, 1964). The single-operator linear or *incremental* and *all-or-none* models have come to be used as standards of a sort (whipping boys might be more accurate) in the field of learning models. Both are extremely simple one-parameter models which are patently incapable of handling the more complex aspects of human verbal learning. However, they do make different predictions about certain features of learning data, and, as it turns out, they lead to quite different optimal presentation strategies.

The *incremental* model is a simplification of the two-operator model previously introduced, in which forgetting is assumed to be a negligible factor; i.e., only permanent learning is taken into account. Specifically, following the n^{th} study trial,

$$q_{i,n+1} = a q_{i,n} \quad (14)$$

which is equivalent to (1), and so from (4)

$$q_{i,n+1} = a^n q_1 \quad (15)$$

where q_1 is the initial error probability, which for present purposes will be assumed to be generally less than 1. Notice that (14) is applied on every study trial, that application of (14) is independent of the student's response on the preceding test, and that the

response gives no information about the state of learning of the item. If the item has received n study trials, then (15) holds, independent of the preceding response history. Note also the implicit assumption that a is the same for all items in the list so that q_n , the *average* probability of an error on n , is also equal to (15).

The *all-or-none* model can be characterized as a two-state Markov process with an absorbing state. It is assumed that each item may be in one of two states: U , an unlearned state in which nothing has been learned about the correct response, and L , a learned state in which the correct response is always given. It may help to think about this model in terms of the utilization of mnemonic codes. On each trial, the student searches through his memory for anything which will help him remember the S-R association. Before such a mnemonic link is discovered the student will be in U . Once a link is found for an item, learning is complete for that item. The model can be described by a transition matrix and column vector of response probabilities.

		State on trial n+1		P(Incorrect Row Vector)
		L	U	
State on trial n	L	1	0	$\begin{bmatrix} 0 \\ q_1 \end{bmatrix}$
	U	c	1-c	

On each study trial, if an item is still in U , then with probability c it is learned, otherwise it remains in U . If an item is in L , it remains there. (As in the incremental model, there is no forgetting.) It is assumed that the parameter c is the same for all items. The guessing parameter q_1 is usually taken to be the reciprocal of the number of available responses.

Suppose that q_n is the average number of items in U on trial n . Then, on the average, a proportion c of the items will change to L

on trial $n + 1$, and $1 - c$ will remain in U . Therefore

$$q_{n+1} = (1 - c)q_n$$

which is identical in form to (14) if $1 - c$ is substituted for a , and hence from (15),

$$q_{n+1} = (1 - c)^n q_1. \quad (16)$$

In short, although the two models are based on different ideas about the psychological processes, the same mean learning curve is predicted by both.

There are some important differences in the predictions of the two models, however. These differences hinge upon the response-dependent character of the all-or-none model. Specifically, an error is an *observable recurrent* event in this model; if the student makes an incorrect response on some item, then the item must have been in U , and none of the preceding study trials had any effect on the likelihood of a correct response for that item. Comparison of response-dependent statistics from the two models points up the difference clearly. Consider the likelihood that an error on trial n is followed by an error $n + 1$, $\Pr(e_{n+1} | e_n)$. According to the incremental model, the error probability on trial $n + 1$ depends only on the fact that there were n previous study trials, so that $\Pr(e_{n+1} | e_n) = a^n q_1$. In the all-or-none model, on the other

hand, the error on n indicates that the item had not been learned up to that point. For an error to occur on the next trial, it must be the case that, with probability $1 - c$, the n th study trial was also ineffective, and that student failed to guess the correct answer with probability q_1 . Putting these two events together, for the all-or-none model, $\Pr(e_{n+1} | e_n) = (1 - c)q_1$. This probability is not a function of the trial number; it is predicted to remain constant throughout learning by the all-or-none model.

Thus, the conditional probability of an error is predicted to remain constant in the all-or-none model and to decrease at an exponential rate by the incremental model. In general, data from verbal learning experiments are closer to the constancy predicted by the all-or-none model than the exponential decrease predicted by the linear model, although a slight decrease is typically observed (Calfee, Atkinson, & Shelton, 1965).

Now we turn to the question of optimal decision strategies based on these two models. The problem will be posed in the following way. Suppose that a fixed number of S-R presentations are to be given, and that the entire sequence of successes and errors for each item in the list is available. How may this history be used to choose an item on each presentation such that learning is optimal? The solution based on the all-or-none model can be understood intuitively as follows. Presentation of learned items is a waste of time, because once an item is learned it remains learned. Hence, the problem becomes one of locating that item which has the best chance of being in the unlearned state. Errors provide this information, in the sense that they identify an item as being in U at the time that the error occurs. The more successes since the last error, the less likely an item is to be in U . *Hence for the all-or-none model, the optimal decision strategy is to present that item with the fewest successes since the last error.*

The technique used to obtain a formal solution of the problem (as opposed to the intuitive argument given above) is *dynamic programming* (Bellman, 1957). Suppose a sequence of D decisions must be made, and it is desirable that the sequence of decisions yield an optimal outcome of some sort. Consider the last or D th decision first. That choice should be made that produces the greatest immediate gain since there are no more decisions to be made. Now take the $(D - 1)$ th decision. The optimal decision at this point is the one which maximizes the gain to be collected from that decision and the last one, assuming that the last decision is optimal. The technique involves working backward in this fashion, always making that decision which yields the highest return from that point on, assuming the remaining decisions are optimal.

In the learning problem, the choice of a particular item constitutes the decision, and the transition of an item to L comprises the gain. (Note that this unobservable gain is paid off at the time of the final test.) The application of dynamic programming is simplified if the following Markovian property of the gain function holds: the payoff arising from the last k decisions depends only on (a) the accumulated gain just prior to the first of these decisions and (b) whatever the last k decisions are. In particular, the payoff from the last k decisions must *not* depend upon the particular manner in which the gain

in (a) was accumulated. Since in our problem the gain is dependent only on performance on the final test trial, the Markov property clearly holds.²

While the gain or payoff is not realized until the final test, it is possible to estimate the probability of gain for each item that might be presented on any trial. This gain is the probability that an item is transferred to L if it is presented. If the item is still in U , then if it is chosen for presentation the likelihood of a transfer is c . If it is in L , then the gain is necessarily zero. Thus we want to find that item with the highest probability of being in U . All events prior to the last error can be completely disregarded, since at that point we know the item was in U and none of the preceding study trials had been effective. It is not difficult to show that the more successes since the last error, the smaller the probability that an item is in U (Appendix B). Hence, the largest payoff on any trial is obtained by presenting that item whose history indicates the fewest successes since the last error.

The next problem is to determine the conditions under which selection of the item which yields the highest (expected) payoff is also an optimal decision strategy. That is, under certain conditions, choice of an item which has the highest payoff on a particular presentation is not necessarily the optimal choice in the long run. It is useful to introduce the concept of *largest immediate gain* (LIG) to refer to the maximum increment in payoff on any single trial, disregarding the potential consequences of the selection at any later stage of the process. If the learning process can be described as a Markov process with observable states, and given some minimal constraints on the gain function, then it can be shown that the *LIG decision*

²As an example of a non-Markovian situation, suppose that the cost of presenting items to the student enters into the evaluation, so that we are trying to maximize amount learned and also minimize the cost of operating the system. Assume also that the cost of presenting item i depends on what item had just been presented, which might be true if the items are messages on a serial device such as a magnetic tape recorder. Substantial cost in time between presentations might be incurred if the items are not close to each other. In this instance the payoff on presentation $n+1$ depends on the payoff on n and the item just presented.

is an optimal decision under dynamic programming (Appendix C).

As mentioned above, for the all-or-none model, the LIG decision is to present the item with the fewest successes since the last error. For the incremental model, the gain associated with item i is the increment in p_i which accrues by presenting i . If an item has previously received n presentations, then its gain will be

$$(1 - a^{n+1}q_1) - (1 - a^nq_1) = a^n(1 - a)q_1. \quad (17)$$

In the incremental model, the gain depends only on how many times an item has been presented and is not a function of the response history. Since (17) decreases as n increases, the LIG decision is to present that item with the fewest presentations. Equivalently, an optimal strategy is to present each item in the list once, then cycle through the list a second time, etc., until the available number of presentations have been used up.

In summary, the optimal decision strategy for the incremental model is to adopt the standard procedure used in most paired-associate studies—present the list of items in random order for as many trials as are available. The all-or-none model implies that a form of *dropout* strategy should be most efficient. An item that has reached a criterion of k successes since the last error should not be presented again until all other items have reached the same criterion.

Research on the relative effectiveness of the dropout procedure is limited. The dropout procedure was apparently introduced half a century ago by Woodworth (1914) but has been only intermittently employed since, mostly as a means of equating the learning of individual items (e.g., Madden, Adams, & Spence, 1950). (Interestingly enough, Woodworth concluded that even when the dropout procedure was used, the more quickly an item reached criterion, the more likely it was to be recalled on subsequent tests.) Battig (1965) argued for the dropout procedure as a remedy for various technical difficulties in the paired-associate method and mentioned that in one experiment the dropout procedure "require[d] significantly fewer errors (as well as less time and fewer total item presentations) to reach a criterion of an errorless trial...[p. 7]." No data were reported.

Dear, Silberman, Eastavan, & Atkinson (1967) have reported two studies designed to evaluate the relative efficiency of the dropout procedure. The stimuli were 32

distinctive 2-digit numbers. The correct response to each stimulus was one of 4 push-button switches. Each response was the correct response for a different set of 8 stimuli, so the task was in some ways analogous to a concept identification problem. The list of stimuli was divided into two sets of 16 items. For each subject, items from one set were presented on alternate trials in the standard fashion, while on the remaining trials, items from the other set were presented using the dropout technique. In the first study, each of 81 subjects was run for 640 anticipation (test-study) presentations, followed by 3 test cycles through all 32 items in the list. (For 44 of these subjects, there were 2 responses instead of 4, but there was no major difference between the groups in the pattern of results.) The number of correct responses on the final tests was slightly higher for dropout items than for standard items, but the difference was not statistically reliable. This was in spite of the fact that a theoretical analysis (assuming that learning was all-or-none) indicated that about 75% of the dropout items should have been learned, but only 45% of the standard items.

In the second study, in which 36 students participated, the only major change was that the training period was terminated when a correct response was given to 10 of the 16 items being presented in a standard manner. Test performance on the standard items was significantly *better* than on the dropout items. From these results, Dear et al. were led to a rather pessimistic conclusion about the efficiency that might be achieved by variation in presentation strategies.

It may be premature to give up hope too quickly, however, especially in light of the other studies referenced above. It would seem that the dropout procedure should have certain advantages over the standard procedure if for no other reason than that certain items are often easier to learn than others. Thus, in learning a foreign language,

less time should be spent on cognates than on unfamiliar items and pseudo-cognates. Dear et al. attempted to select a set of stimuli-response pairs of about equal difficulty, and so the comment above may not be especially germane to their results.

The two models used for optimization are both incomplete in one obvious respect, viz., neither incorporates a mechanism to handle forgetting, which arises from the efforts of the student to learn other items in the list. Models exist which provide representation for both learning and forgetting processes (Atkinson & Craighero, 1964; Calfee & Atkinson, 1965). However, it has proven difficult to derive optimal strategies based on these more complex models.

As an alternative to the use of a more complicated model, one might change the experimental procedure in certain ways. The anticipation technique used by Dear et al. has the undesirable feature that once an item reaches some performance criterion in the dropout procedure, it is not likely to be presented for study for some time, and consequently no tests are run for that period. If for any reason the item is forgotten, the teaching system remains insensitive to the loss. A more appropriate experimental procedure might involve separation of test and study events. Every item would be tested at regular intervals, but the choice of items for study would be determined by an optimal decision strategy (also, cf. Battig, 1965). As an alternative, the anticipation technique might be used but the teaching system might keep a running tally for each item of how many trials have passed since it was last presented. After some maximum was reached, an item would be presented for review. Both of these techniques would produce a better history on which to base decisions, a history which would be less likely to be misleading because of short-term memory and chance guessing effects.

IV EFFICIENCY OF DECISION STRATEGIES

Calculation of the theoretical efficiency of various decision strategies has proven rather difficult. The application of dynamic programming, though it allows one to find an optimal strategy, does not generally produce a closed-form expression for the gain in efficiency of that strategy over any alternative that might be proposed. In order to get an idea of the relative merits of the dropout procedure, a series of computer simulations were carried out. In each simulation run, Monte Carlo data were generated for 100 stat-students, each learning a 10-item list. Independence of items was assumed, and the initial guessing rate q_1 was set at .1. An anticipation technique was used, and to get the process started there were two preliminary trials or cycles through the entire list. N additional presentations were then allotted to various items according to some particular decision strategy. The efficiency of the strategy was determined after the N presentations by computing the average number of items in L for the all-or-none model, or the mean of q_i for the incremental model.

Three decision strategies were evaluated: (a) the dropout procedure, (b) the standard trial procedure, and (c) a strategy based on an intuition that difficult items would have relatively more errors, hence items should be presented relatively in proportion to the error rate. This last strategy which appeared to be reasonable on intuitive grounds was included for comparison with the solution derived from a more formal approach, dynamic programming.

The first series of simulation studies investigated the efficiency of the dropout procedure compared with the standard procedure, assuming an all-or-none model and identical learning rate parameters for all items. Strategy (c) was also included for comparison. The results are shown in Table 1 for values of c from .3 to .025, a range typical of the results of paired-associate experiments, and for N equal to 50, 80, and 160.

Table 1. Proportion of Items in L at End of Training Following Standard, Dropout, and Error-Proportion Decision Strategies (Simulation Data Based on All-or-None Model, Homogeneous Parameters)

N	c	Decision Strategy		
		Standard	Dropout	Error-Proportion
50	.3	.91	1.00	.97
	.1	.53	.65	.56
	.05	.33	.31	.30
	.025	.16	.17	.16
80	.3	.98	1.00	1.00
	.1	.62	.83	.70
	.05	.36	.48	.44
	.025	.22	.26	.24
160	.3	1.00	1.00	1.00
	.1	.85	1.00	.96
	.05	.61	.80	.69
	.025	.38	.46	.36

These data, while limited, provide some idea of the relative efficiency of the various procedures. When there is a difference, the dropout procedure is most efficient, the standard procedure least efficient; the strategy based on error proportions falls in between. It appears that the efficiency of the dropout procedure depends on both c and N . For certain combinations of these variables, the dropout procedure is considerably more efficient (e.g., $c = .1$, $N = 80$). However, if N is large *and* if learning is fast, then all the items will be in L at the end of training; conversely, if the number of available presentations is small relative to the learning rate then not many items will have been learned, regardless of the presentation procedure. For certain intermediate conditions which

cannot be described in a simple fashion, substantially more items will be learned if the dropout procedure is used.

One slight change in procedure provides a more certain estimate of efficiency. The assumption of a fixed number of presentations facilitated the application of dynamic programming to the optimization problem. However, suppose that instead of having a fixed number of presentations, the session is terminated only after some proportion of items in the list have reached a predetermined performance criterion (so many successes in a row). The performance criterion should reflect the chance guessing rate—the higher the probability of being correct by chance, the more stringent should be the criterion.

Again, assume that learning proceeds in an all-or-none fashion. The discussion will be simplified if q_1 is assumed to be 1; the general result can be obtained in closed form, but is less informative than the result below, which will depart little from the general result if q_1 is reasonably close to 1. Suppose the proportion of items which must be learned (i.e., reach the performance criterion) is set at some value, x .³ Taking the standard procedure first, on the average we can expect some proportion of the items to be learned on trial 1, an additional proportion on 2, etc. At some critical trial, the proportion of items learned will equal x . For the standard procedure, let k_S be the mean number of presentations required before the last error.⁴ The

³It should be noted that two criteria have been mentioned, (a) a performance criterion for each item of some number of consecutive successes, and (b) a list criterion for terminating the training session. Criterion (b) is attained when a proportion x of the items in the list have reached the performance criterion (a).

⁴For convenience in deriving results for this section, we are considering the number of presentations to the last error, not the number of presentations required to complete the criterion run. Since these criterion trials add a constant to the numerator and denominator of the efficiency ratio, k_S/k_D , an efficiency ratio based on trials to criterion will be attenuated to an extent that depends on the length of the performance criterion and the actual values of k_S and k_D .

distribution of E , the presentation of last error is

$$\Pr(E = n) = (1 - c)^{n-1} c. \quad (18)$$

For the last error to occur on the n^{th} presentation, it must be the case that learning did not occur on $n - 1$ previous opportunities, but then occurred on the n^{th} presentation. The problem is to find k_S such that

$$x = \sum_{i=1}^{k_S} \Pr(E = i) = \sum_{i=1}^{k_S} c(1 - c)^{i-1} \quad (19a)$$

$$= 1 - (1 - c)^{k_S}$$

Rearranging terms and taking logarithms to base 10,

$$(1 - c)^{k_S} = 1 - x,$$

$$k_S = \frac{\log(1 - x)}{\log(1 - c)}. \quad (19b)$$

In the dropout procedure, an item is presented only until it reaches criterion. Thus, on the average a proportion of the items equal to $\Pr(E = 1)$ will be presented once, a proportion $\Pr(E = 2)$ twice, and finally the last item reaching the list criterion will require k_S presentations. The average number of presentations for that portion of the list which reaches criterion will therefore equal

$$\sum_{i=1}^{k_S} i \Pr(E = i) = \sum_{i=1}^{k_S} i c(1 - c)^{i-1}$$

$$= 1 - \frac{(1 - c)^{k_S} (1 + k_S c)}{c} = \frac{1 - (1 - x)(1 + k_S c)}{c}. \quad (20)$$

It is not exactly clear, given completely random selection of those items remaining to be learned in the dropout procedure, just what the distribution of the number of presentations would be for the unlearned items which remain when the list criterion is finally reached. We should not be too far off, however, if we assume that the noncriterion items will have received about as many trials as the last items to be learned, viz., k_S trials. Finally, then, for a proportion x of the items in the

dropout procedure the mean number of presentations will be given by (20), and for the remaining $1 - x$ of the items the number of presentations will be approximately k_S . Combining these results, the mean number of presentations per item using the dropout procedure, k_D , will be

$$k_D = \frac{1 - (1 - x)(1 + k_S c)}{c} + (1 - x)k_S$$

$$= \frac{x}{c} \quad (21)$$

For that range of c which is of interest, $.3 > c > .01$, $\log_{10}(1 - c)$ is very nearly equal to $-c/2$. From (19b), substitution of this approximation for $\log(1 - c)$ yields

$$k_S = \frac{-2 \log(1 - x)}{c} \quad (22a)$$

from which is obtained the final result,

$$\frac{k_D}{k_S} = \frac{x}{-2 \log(1 - x)} \quad (22b)$$

In short, the relative efficiency of the dropout procedure depends only on x , the proportion of items which must reach the performance criterion for a session to be terminated, and does not depend on the learning rate within the boundary values mentioned. From (22b), the dropout procedure requires 57% as many presentations as the standard procedure for $x = .8$; for a more stringent list criterion, $x = .95$, the dropout procedure requires only 25% as many presentations.

The argument above strictly holds only if the list length is infinitely large. For lists of finite size, the expected value of k_S obtained by (19a) may not be exactly correct. (Estes, 1959, pp. 36 ff., gives the distribution function for k_S for finite lists, but the expected value involves the incomplete sum of a binomial distribution.) The approximation in (22b) was evaluated by computer simulation for 10-item lists and for various values of x and c . Agreement between the obtained and predicted efficiency ratios was good except for $x = 1.0$, in which case the ratio appeared to be approaching a value of about .4. A reasonable conjecture would seem to be that (22b) holds up for short lists for all criteria up to $(M-1)/M$, but not for $x = 1$. Simulations also showed that for $q_1 = .9$, the approximation from (22b) was quite accurate.

In summary, given a fixed number of presentations, the dropout procedure was shown by application of dynamic programming to be optimal. The relative efficiency of the dropout procedure compared to the standard procedure varied in a complex fashion with the learning rate and the number of presentations. Under certain conditions, final test performance was measurably higher under the dropout procedure, but other times the difference in procedures was quite small. By requiring that a session last until a criterion proportion of items has been learned, the relative efficiency of the dropout procedure was determined to a good approximation and was found to depend only on the criterion. If the list criterion was 90% or higher, the dropout procedure should require less than half as many presentations as the standard procedure—a substantial gain in efficiency.

V HETEROGENEITY OF LEARNING RATE PARAMETERS

If one assumes that some items are easier to learn than others, then the dropout procedure might prove efficient because it provides more practice on difficult pairs. It seemed worth investigating this possibility in the context of both models. Assume that a list with two items is to be learned, and that item 1 is much easier to learn than item 2. In the all-or-none model, this is equivalent to $c_1 > c_2$.

The dynamic programming solution handles this case without any need to estimate c_1 and c_2 ; the easier item will go to L relatively early, and the majority of the presentations will be given to item 2.⁵

Next consider the incremental model, where $a_1 < a_2$. If the two parameters are

⁵If the parameters c_1 and c_2 were known prior to the experiment, the optimal decision strategy would actually be more complicated. The nature of this complication may be seen by considering the decision rule if the last response to both items was an error. If $c_1 = c_2$, there is no optimal decision. If $c_1 > c_2$, then item 1 should be presented, since the expected gain, c_1 , is greater than for item 2.

More generally, the *a posteriori* probability (based on successes since last error) that each item is in U is not sufficient to make the decision, but must be weighted by the learning rate parameter. The strategy of the dynamic programming solution in this instance is to teach the easy items first. In most instances, however, no *a priori* estimate of the learning rate parameters will be available, unless the assumption is made that the relative difficulty of an item can be estimated over a group of students. Smallwood (1962) considers this possibility in a different context.

known *a priori*, then the dynamic programming solution can be readily computed from the parameter values and the number of times each item has been presented, n_1 and n_2 .

From (18), the LIG for item i will be $a_i^{n_i}(1 - a_i)$.

Another approach which will be useful later involves finding the optimal proportion of the available presentations which should be assigned to item i . Assuming known parameters and letting N equal the available presentations, for the 2-item case we can write

$$E(q) = a_1^k + a_2^{N-k}. \quad (23)$$

Differentiating (23) with respect to k and solving for the minimum yields the following result for the optimal proportion of presentations to be assigned to item 1,

$$z_{\text{opt}} = \frac{1}{\log a_1 + \log a_2} \left\{ \log a_2 + \frac{1}{N} \log \left[\frac{\log a_1}{\log a_2} \right] \right\}. \quad (24)$$

The problem is more difficult when, as is generally true, *a priori* estimates of the learning rates do not exist. A moments estimate of a_i can be readily obtained from $T_{i,n}$, the expected total errors for item i in n presentations of the item, since

$$T_{i,n} = q_i \sum_{j=0}^{n-1} a_i^j = q_i \left[\frac{1 - a_i^n}{1 - a_i} \right] \quad (25a)$$

implying that

$$a_i = q_i \left[\frac{T_{i,n} - 1}{T_{i,n} - a_i^n} \right] \approx q_i \left[\frac{T_{i,n} - 1}{T_{i,n}} \right] \quad (25b)$$

As it turns out there are some practical problems in using (25b) for optimization. The

term a_i^n in the denominator will be close to zero for reasonably large values of n . If n is small and a_i is large (fast learning), then (25b) leads to an estimate of a_i biased toward 0, so that learning may be slower than estimated. An additional problem is that $T_{i,n}$ is an *expected* value over many items, while the available datum for estimation in an optimization schedule must be based on a single item. Accordingly, an estimate of a_i based on (25b) might be quite bad. It appears that the worst case arises if by chance a student guesses correctly during the first one or two presentations of a difficult item. Then a_i will be estimated to be 0, and the item is unlikely to be presented again. A practical answer to this problem has already been suggested, viz., schedule tests on a regular and independent basis.

Another solution is to apply the dynamic programming solution probabilistically rather than deterministically. Recall that the optimal decision strategy was to choose that item for which $a_i^{n_i}(1-a_i)$ is greatest. Suppose the gain, relative to the rest of the items

$$g_i = a_i^{n_i}(1-a_i) / \sum_{j=1}^M a_j^{n_j}(1-a_j) \quad (26)$$

is computed. If a_i is estimated as 0, an arbitrary minimum value (e.g., a constant or $\frac{1}{n_i}$)

is substituted for a_i . On each presentation, the probability that i selected will be g_i .

The probabilistic choice means that even those items that have been tentatively identified as easy will be presented occasionally, rather than being "trapped."

A second series of simulation runs were carried out to evaluate various decision strategies, assuming either the all-or-none or incremental model, and using heterogeneous parameter sets. The three decision strategies previously described were investigated, together with a fourth strategy based on the probabilistic solution just described based on running estimates for a_i for each item.

Both the incremental and all-or-none models were simulated. Five items in the list were assigned a learning rate parameter of c_1 (or $1-a_1$), and the other five were assigned a

parameter of c_2 . The number of available presentations was set at 80. The entries under "Parameter Known" give the expected state values if the available presentations were distributed optimally using (24).

The results of these studies are presented in Table 2. For the all-or-none model, the dropout procedure is best; the standard procedure, worst; and the other two strategies fall between. As was suggested earlier, the dropout procedure handles parameter heterogeneity with no complications.

For the incremental model, two questions seem pertinent: (a) When heterogeneity of learning rates exists, how much gain improvement is possible by optimal scheduling based on *a priori* estimates of the parameters? and (b) When the parameters are unknown, how close do various presentation strategies come to achieving the potential gain in efficiency? The answer to the first question, at least for the parameter sets in Table 2, is that the potential gain is relatively small. Only when some items are learned extremely fast—one or two trials—and other items quite slowly does there seem to be any substantial difference between strategies. As to the second question, only for the most extreme case ($a_1 = .4$, $a_2 = .95$) do the various decision strategies seem equally effective. In particular, the dropout procedure seems as efficient as the more elaborate procedure based on the use of running estimates of the a_i to calculate the immediate gain for each item.

It may be helpful to summarize the major points in this section on dynamic strategies. The derivation of optimal presentation techniques based on backward induction was introduced, and it was shown that for certain criteria of optimality and a general class of learning models, presenting that item which yielded the maximum immediate gain was an optimal strategy. From this result, it was possible to show that a dropout procedure was optimal for an all-or-none learning model, whether the learning rate was the same for all items in the list or not. If an incremental model was assumed, then the standard procedure was optimal if all items were learned at the same rate; otherwise, a more complex decision rule was required based on running estimates of the learning rates and calculation of largest immediate gain. Simulation studies suggested that for the case of heterogeneous learning rates the dropout procedure performed as well as the more complex strategy above.

Table 2. Mean State Values at End of 100 Presentations, 10-Item List for Various Decision Strategies, All-or-None and Incremental Models with Heterogeneity of Parameters

Model		Decision Strategy				Known Parameters
All-or-none		Standard	Dropout	Running Estimate	Error Proportion	
c_1	c_2					
.60	.05	.31	.14	.23	.20	-
.30	.10	.20	.01	.12	.11	-
.40	.05	.29	.18	.28	.26	-
.15	.05	.40	.28	.32	.38	-
.20	.025	.41	.36	.42	.40	-
Incremental						
a_1	a_2					
.40	.95	.30	.25	.26	.26	.24
.70	.90	.19	.19	.19	.19	.15
.60	.95	.30	.29	.29	.29	.26
.85	.95	.40	.44	.41	.41	.40
.80	.97	.44	.50	.45	.46	.41

In general, only slight gains in efficiency were achieved by taking into account heterogeneity in learning rates, assuming an incremental model.

Under certain conditions, the gain in performance from the dropout technique for a fixed number of presentations was substantial. However, in many other cases, the gain was negligible. More consistent outcomes were found by assuming that the training session was continued until some proportion x of the items had reached a performance criterion. The relative efficiency of the dropout procedure under these conditions depended only x , and for reasonable values of x the gain was considerable.

In this connection, it might be useful to point out some desirable features for experiments designed to evaluate the effectiveness of these two procedures. First, as just mentioned, teaching should continue to a list criterion. Second, the probability of chance successes should be relatively small, since otherwise the Bayesian estimate of the probability that an item has been learned is not very reliable. Third, each item should be tested sufficiently often so that correct responses arising from short-term retention do not carry much weight in the decision of the teaching system.

VI CONCLUSIONS AND IMPLICATIONS

The theoretical analyses of optimality discussed above have centered around a stimulus-response procedure of teaching—a conceptually simple paradigm and one of the few for which reasonably adequate models have been developed. A wide class of instructional problems can be expressed in this framework. More importantly, the theoretical models considered span the range of assumptions about how learning occurs from the "rote-learning" incremental model (representative of the Thorndikian establishment of connections by sheer repetition) to the "insight" all-or-none model based on active search for mnemonic links. The models considered are oversimplifications, but in various situations they provide good approximations to the learning process, approximations which are mathematically tractable. Because mathematical analysis rapidly becomes more difficult as one complicates a model, it would seem a reasonable strategy to press these simple representations as far as possible. It is my suspicion that from the standpoint of optimization the most serious problem with these models is that no provision is made for interitem interference—for the fact that in learning one item, the student may form an association that makes it difficult to retrieve

other associations that were previously formed. This interference becomes more important with increased similarity among items in a list. In practical applications, as list length increases the amount of interitem similarity also tends to increase.

It is apparent that the psychological processes used by a student in learning associations between discrete stimuli and responses—no matter how complex either the stimulus or response terms may be—represent only a fraction of the cognitive skills in the student's repertoire. Very likely, different cognitive processes are operative when the student is asked to learn conceptual relations (e.g., physics), remember a body of facts in an organized or structural fashion (e.g., American history), or become skillful at applying transformations of various sorts (e.g., algebra). The examples are for illustration only. In any curricular area such as reading or arithmetic, many different cognitive skills are needed. A promising approach to achieving instructional efficiency would be to partition an area according to the requisite skills and seek to optimize the instructional system with respect to the development of those specific skills, as in the examples in this paper.

APPENDIX A EVALUATION OF BLOCK-SIZE FUNCTION

The following proof was suggested by J. C. Falmagne. Define the two functions

$$\phi_N(k) = \frac{1 - b^{kN} v^N}{1 - b^k v}, \quad \psi_N(k) = \frac{1 - b^{kN}}{1 - b^k} \quad (A1)$$

The last two terms in (8) are equal to $\phi_N(k)/\psi_N(k)$, and the problem is to show that $v < 1$ implies $E(q)$ decreases monotonically with k , and $v > 1$ implies $E(q)$ increases with k . Equivalently, it must be shown that for all ℓ

$$v < 1, \quad \frac{\phi_N(k + \ell)}{\psi_N(k + \ell)} \geq \frac{\phi_N(k)}{\psi_N(k)} \quad (A2)$$

$$v > 1, \quad \frac{\phi_N(k + \ell)}{\psi_N(k + \ell)} \leq \frac{\phi_N(k)}{\psi_N(k)}.$$

Notice that (A2) holds if $N = 1$, since $\phi_1(k) = \psi_1(k) = 1$ for all k . We show that (A2) is true for all N by induction for $v < 1$; the proof for $v > 1$ is similar.

Suppose that (A2) holds for some N . It must be shown to hold for $N + 1$, viz.,

$$\frac{\phi_{N+1}(k + \ell)}{\psi_{N+1}(k + \ell)} \geq \frac{\phi_{N+1}(k)}{\psi_{N+1}(k)} \quad (A3)$$

Both ϕ and ψ can be expressed in series form,

$$\phi_N(k) = \sum_{i=0}^{N-1} b^{ki} v^i, \quad \psi_N(k) = \sum_{i=0}^{N-1} b^{ki} \quad (A4)$$

so that (A3) is equivalent to

$$\frac{\phi_N(k + \ell) + b^{(k+\ell)N} v^N}{\psi_N(k + \ell) + b^{(k+\ell)N}} \geq \frac{\phi_N(k) + b^{kN} v^N}{\psi_N(k) + b^{kN}} \quad (A5)$$

Crossmultiplying and cancelling common terms, (A5) becomes

$$\begin{aligned} \phi_N(k + \ell) \psi_N(k) + b^{kN} \phi_N(k + \ell) + b^{(k+\ell)N} v^N \psi_N(k) \geq \\ \phi_N(k) \psi_N(k + \ell) + b^{(k+\ell)N} \phi_N(k) + b^{kN} v^N \psi_N(k + \ell). \end{aligned} \quad (A6)$$

From the inductive hypothesis,

$$\frac{\phi_N(k+\ell)}{\psi_N(k+\ell)} \geq \frac{\phi_N(k)}{\psi_N(k)} \quad \text{implies} \quad \phi_N(k+\ell) \psi_N(k) \geq \phi_N(k) \psi_N(k+\ell) \quad (\text{A7})$$

hence the crossproducts of ϕ and ψ may be eliminated without prejudice to the hypothesis. Rearranging terms and expressing ϕ and ψ as series, (A6) will be true if

$$\phi_N(k+\ell) - b^{\ell N} \phi_N(k) \geq v^N \{ \psi_N(k+\ell) - b^{\ell N} \psi_N(k) \} \quad (\text{A8})$$

or

$$\sum_{i=0}^{N-1} v^i \{ b^{(k+\ell)i} - b^{\ell N+ki} \} \geq \sum_{i=0}^{N-1} v^N \{ b^{(k+\ell)i} - b^{\ell N+ki} \}. \quad (\text{A9})$$

However, $b \leq 1$, $i < N$, and so $b^{ki} \{ b^{\ell i} - b^{\ell N} \} \geq 0$. Moreover, $v < 1$, hence $v^i > v^N$ for each i , and thus for each pair of terms in the series the inequality is true.

APPENDIX B
DERIVATION OF A POSTERIORI PROBABILITY
OF STATE L IN ALL-OR-NONE MODEL

Let K be the occurrence of k successes since the last error, U be the event that the item is in U after K , and let X represent the relevant features of the learning history, viz., the last error and the $k+1$ study trials (one following the last error, k following the k successes). Then the task is to find $\Pr(U|K \text{ and } X)$, which may be rewritten by Bayes formula as

$$\Pr(U|K \text{ and } X) = \frac{\Pr(K|U \text{ and } X)P(U|X)}{P(K|X)} \quad (B1)$$

Considering the terms one at a time, $\Pr(K|U \text{ and } X)$, since it is conditional on U , is the likelihood that k successes occur by chance, $(1-q_1)^k$. The second term, $\Pr(U|X)$, is $(1-c)^{k+1}$, since the student must fail to learn the item on each of the $k+1$ opportunities. The denominator, $\Pr(K|X)$, must be broken into two parts, depending on whether the item was learned during one of the $k+1$ study trials, or not learned at all. Accordingly

$$\begin{aligned} \Pr(K|X) &= \Pr(K|U \text{ and } X)\Pr(U|X) + \sum_{i=1}^{k+1} \Pr(K|L \text{ on } i \text{ and } X)\Pr(L \text{ on } i|X) \\ &= (1-q_1)^k (1-c)^{k+1} + \sum_{i=1}^{k+1} c\{(1-c)(1-q_1)\}^{i-1} \\ &= (1-q_1)^k (1-c)^{k+1} + \frac{c[1 - \{(1-c)(1-q_1)\}^{k+1}]}{1 - \{(1-c)(1-q_1)\}}. \end{aligned} \quad (B2)$$

Further simplification yields the final result,

$$\Pr(U|K \text{ and } X) = \frac{\{(1-c)(1-q_1)\}^{k+1} \{1 - (1-c)(1-q_1)\}}{c(1-q_1) + q_1\{(1-c)(1-q_1)\}^{k+1}}. \quad (B3)$$

The derivative of (B3) with regard to k is negative, which means that $\Pr(U|K \text{ and } X)$ decreases uniformly as k increases. In other words, the item with the fewest successes since the last error has the highest probability of being in U .

It is obvious that if $q_1 = 1$, then $\Pr(U|K \text{ and } X) = 0$ for $k \geq 1$; i.e., if the guessing rate is zero, a single success is sufficient to imply that learning has occurred. Inspection of (B3) suggests that as q_1 decreases to .5, $\Pr(U|K \text{ and } X)$ will increase. Thus, the higher the guessing rate, the less certain one can be that the process is in L for any fixed value of k .

APPENDIX C CONDITIONS FOR OPTIMALITY OF LARGEST-IMMEDIATE-GAIN STRATEGY

Suppose that a student is to be taught a list of M stimulus-response pairs and only N presentations are available. Prior to each presentation, the instructor chooses a pair from the list to present to the student. The optimization problem of concern here involves finding a *decision* procedure for selecting items such that expected number of correct responses on a final test (or some monotonic function thereof) is maximized. The performance level on the final test will be referred to as the *payoff*.

The optimum decision strategy is a function of the *history* and *gain function* for each pair. A *history*, for purposes of this paper, is the sequence of correct and incorrect responses for an item up to some point in the presentation sequence. The *gain function* defines the expected increment in payoff which accrues if a particular item is presented. In the present discussion, the payoff is measured by the number of correct responses on the final test. Accordingly, the gain function will be some function of the increase in the probability of a correct response obtained by presenting an item.

The relation between payoff and gain may be further clarified by looking at the decision problem from a different angle. Suppose that k presentations remain to be made at some point. This portion of the presentation sequence can be represented by a tree diagram with k levels, laying out the M^k different sequences which might occur. Given the history at k and specification of a learning model, the expected payoff for each sequence may be calculated. One solution of the optimization problem is to calculate all possible outcomes prior to each decision, and present the item which leads to the highest payoff. This approach is obviously unworkable unless M and k are both quite small. A major concern is to discover simpler and more efficient means than enumeration of arriving at optimal decisions.

It will be recalled that gain is defined as the expected increment in payoff obtained by presentation of an item at some point in the sequence. It is not generally true that selection of that item which produces the largest expected gain on a presentation is an optimal choice in terms of the expected payoff. As a simple counterexample, suppose a list consists of two items, A and B, and there are two presentations remaining. Suppose that the probabilities of a correct response following one and two presentations for A are .20 and .40, respectively, and for B are .10 and .60. Suppose further that the payoff is the probability of a correct response following the two presentations. The possible outcomes can be represented by a tree diagram:

First Decision	Expected Gain	Second Decision	Expected Gain	Final Payoff
<div> <div></div> <div>A</div> <div></div> </div>	.20	A	.20	.40
		B	.10	.30
<div> <div></div> <div>B</div> <div></div> </div>	.10	A	.20	.30
		B	.50	.60

By this enumeration it can be seen that the optimal strategy is to present B twice in succession, whereas if gain were being maximized, one would present A twice in succession. A major purpose of this section is to show that under certain conditions maximization of immediate gain is also an optimal decision.

Assume that the learning process for each item can be expressed as a Markov process with observable states S_1, \dots, S_s . Let the gain function, $g_x(i, n)$, be the increment in the expected payoff of item x in S_i when it is presented on n ; $g_x(i, n+1|n)$ is the gain from presentation on $n+1$, given item x was also presented on n and was in S_i at that point. (If the anticipation procedure is used in the learning task, it is necessary to be quite explicit. The *state* of the item will mean its state prior to the test portion of the presentation.)

The notation is unavoidably complicated by the need at various times to take note of the particular item, the state of the item, and the position in the presentation sequence. Where it is of no consequence, the item subscript will be dropped. The index n is especially troublesome. It may be used to indicate either of two related numbers: (a) position in the presentation sequence counting forward from the first decision, denoted by n , or (b) position in the sequence counting backward from the last decision, denoted by k . A largest-immediate-gain (LIG) decision is optimal under dynamic programming if for all items (C1) $g(i, n) \geq g(j, n)$ for $i \leq j$, and (C2) if $g(i, n) \geq g(i, n+1|n)$. The first condition requires that the states be arranged so that the expected gain decreases monotonically with the state index. For the second condition to be met, if an item in S_i is presented on n and also on $n+1$, the gain on $n+1$ must not be greater than the gain on n . If these conditions are met, the LIG decision is the choice of an item with the smallest current state value.

The proof proceeds by induction, and follows a discussion in Matheson (1964, pp. 22-24) with changed notation. It is obvious that the LIG decision is optimal for the last decision in the sequence, because no further decisions remain to be made. For the inductive hypothesis, assume that with k decisions remaining, the LIG decision is optimal. Suppose that with $k+1$ decisions remaining, item x yields LIG, and item y is an optimal choice. The proof involves showing that on $k+1$, x is at least as good a choice as y .

Recall that $g_x(i, k)$ is the expected gain from presenting item x in state i with k decisions remaining. By assumption $g_x(i, k+1) \geq g_y(j, k+1)$. Moreover, by C2 $g_y(j, k|k+1) \geq g_y(j, k)$, and so

$$g_x(i, k+1) \geq g_y(j, k+1) \geq g_y(j, k|k+1). \quad (1)$$

If y is selected for presentation at $k+1$ as the optimal choice, then $g_z(m, k) = g_z(m, k+1)$ for all $z \neq y$. In other words, there is no change in state and therefore no change in gain from $k+1$ to k except for the item presented at $k+1$. In particular, $g_x(i, k) = g_x(i, k+1)$. Accordingly, x is LIG at decision k , and by the inductive hypothesis is the optimal choice. Given the usual assumption that learning proceeds independently for each item, it makes no difference in what order x and y are presented. The sequences $\{x \text{ on } k+1, y \text{ on } k \text{ and } y \text{ on } k+1, x \text{ on } k\}$ are therefore equivalent so that x must also be optimal on $k+1$, which completes the proof.

The critical role of the gain function is apparent from the work of Karush and Dear (1966a) who showed that for arbitrary gain functions, a choice based on LIG may be far from optimal, and that more complex strategies based on gain accrued over two or more succeeding presentations may be necessary.

Both the incremental and all-or-none models can be described by gain functions satisfying C1 and C2 (cf. Atkinson and Estes, 1963, for details on these learning models). In the incremental model, the states are indexed by the number

of previous study presentations. The probability of a success following the i th study trial is $p_n = 1 - qa^{i-1}$ where q is the initial error probability and a is the learning rate parameter. A suitable gain function is the difference between p_i and p_{i+1} , $g(i,n) = q(1-a)a^{i-1}$, which clearly satisfies both C1 and C2.

It requires a bit more work to handle the all-or-none model, since it must be rearranged in the form of an identifiable state model and a suitable gain function found. The underlying model is a two-state Markov process. In the unlearned state, the probability of an error is q ; on each trial, there is a probability c of transition to the learned state. The learned state is an absorbing state, in which the probability of an error is 0. An appropriate Markov description is obtained by letting the states be indexed in the number of successes since the last error. For derivational purposes, it is easiest to work with Q_i , the probability that an item is in the unlearned state given i successes since the last error, which is (Appendix B)

$$\frac{\{(1-c)(1-q)\}^{i+1} \{1 - (1-c)(1-q)\}}{c(1-q) + q \{(1-c)(1-q)\}^{i+1}} \quad (2)$$

The all-or-none model is a recurrent process, and Q_i is not dependent on either the presentation number n or how many times an item has been presented. Moreover, it is simple to show that Q_i decreases monotonically with i , i.e., that $Q_i - Q_{i+1} \geq 0$.

The gain function $g(i,n)$ will be defined as the (expected) decrement in the probability that an item is unlearned, given that the item is in S_i just prior to being chosen for a test-study sequence on the n th presentation. The decrement depends on the outcome of the test. With probability qQ_i , an error will occur, and the item will be identified as being unlearned with probability 1. The effect of the study event is to reduce this probability to $1-c$, and accordingly the gain from study immediately following an error is c . With probability $1 - qQ_i$, a success will be observed, the *a posteriori* probability will change to Q_{i+1} , and the gain from the study event will be cQ_{i+1} . The anticipation procedure necessitates careful consideration of the sequencing of events. In particular, the gain depends on the outcome of the test portion of a trial and the immediately following study portion, as does the determination of the state S_i of an item. The relation between the events of a trial, the gain associated with those events, and the various indices is shown on the next page for an item which is presented for two trials in succession.

The expected immediate gain from presentation of an item in state i , $g(i,n)$ is just


$$g(i,n) = c[(1 - qQ_i)Q_{i+1} + qQ_i]. \quad (3)$$

Straightforward algebra shows that (3) meets C1, i.e., that $g(i,n) \geq g(i+1,n)$:

$$\begin{aligned} (1 - qQ_i)Q_{i+1} + qQ_i &\geq (1 - qQ_{i+1})Q_{i+2} + qQ_{i+1}, \\ (1 - qQ_{i+1})(Q_i - Q_{i+2}) &\geq (1 - q)(Q_i - Q_{i+1}). \end{aligned} \quad (4)$$

But clearly, $(1 - qQ_{i+1}) \geq (1 - q)$, and since Q_i decreases with i , $(Q_i - Q_{i+2}) \geq (Q_i - Q_{i+1})$.

Presentation n				Presentation n+1			
Test		Probability Unlearned	Study	Test		Probability Unlearned	Study
			State				State
			Gain				Gain



Demonstrating that (3) satisfies C2 requires a little more care. From the two-trial sequence presented above, it can be seen that the problem is to show that

$$(1 - qQ_i)Q_{i+1} + qQ_i \geq (1 - qQ_i)[(1 - qQ_{i+1})Q_{i+2} + qQ_{i+1}] + qQ_i[(1 - qQ_0)Q_i + qQ_0]. \quad (5)$$

The rightmost terms involving qQ_i are easily shown to satisfy the inequality. The left hand terms in $(1 - qQ_i)$ must be changed to a form such as

$$\frac{Q_{i+1}}{1 - qQ_{i+1}} \geq \frac{Q_{i+2}}{1 - q},$$

and the appropriate substitutions made from (2) in order to see that the result is equivalent to $c(1 - q) \geq 0$, which clearly holds.

Thus for the all-or-none model, there is an observable state representation and a gain function satisfying C1 and C2 which are appropriate for the adoption of an LIG strategy. It might be mentioned that there exist other plausible gain functions for this same representation of the model which do *not* satisfy C1 or C2. For example, in the initial stages of working on this problem, the gain was defined as the expected decrement in Q_i , the *a posteriori* probability of being in the unlearned state given i successes since the last error. This gain function does not satisfy either C1 or C2, and hence the LIG strategy is inappropriate. However, on reflection it is apparent that the function is not appropriate for another, more pertinent reason. The gain represented by this function relates to the increase in the *teacher's certainty* that learning has occurred, not in the probability that learning actually takes place. For example, suppose that q_1 is .5, and that c is quite small. The likelihood of chance success is so high that strings of up to three successes are not very diagnostic of the learning state; in other words, Q_i decreases slowly for $0 \leq i \leq 3$. Longer strings quickly become improbable under the hypothesis that the item is unlearned, and so Q_i decreases rapidly. Under these conditions of the parameters and for this particular gain function, the optimal choice is to choose an item with a string of two or three successes—if another success occurs, there will be a big jump in the *a posteriori* probability that the item is learned. (One might think of this gain function as suitable for the type of individual who turns to the back of a mystery novel to find out "who done it?") In summary, the demands on the model in this approach are relatively minimal, but the selection of a proper gain function is critical.

REFERENCES

- Atkinson, R. C., Bower, G. H., & Crothers, E. J. *Introduction to mathematical learning theory*. New York: Wiley, 1965.
- Atkinson, R. C., & Crothers, E. J. A comparison of paired-associate learning models having different acquisition and retention axioms. *Journal of Mathematical Psychology*, 1964, 1, 285-315.
- Battig, W. F. Procedural problems in paired-associate learning research. *Psychonomic Monograph Supplement*, 1965, 1, No. 1.
- Bellman, R. *Dynamic programming*. Princeton: Princeton University Press, 1957.
- Calfee, R. C., & Atkinson, R. C. Paired-associate models and the effects of list length. *Journal of Mathematical Psychology*, 1965, 2, 254-265.
- Calfee, R. C., Atkinson, R. C., & Shelton, T., Jr. Mathematical models for verbal learning. In N. Weiner and J. P. Schade (Eds.), *Progress in Brain Research*, Vol. 17, *Cybernetics of the Nervous System*. Amsterdam: Elsevier, 1965. Pp. 333-349.
- Crothers, E. J., and Suppes, P. *Experiments in second-language learning*. New York: Academic Press, 1967.
- Dear, R. E., Silberman, H. F., Estavan, D. P., & Atkinson, R. C. An optimal strategy for the presentation of paired-associate items. *Behavioral Science*, 1967, 12, 1-13.
- Estes, W. K. Component and pattern models with Markovian interpretations. In R. R. Bush and W. K. Estes (Eds.), *Studies in mathematical learning theory*. Stanford: Stanford University Press, 1959. Pp. 9-52.
- Goldberg, S. *Introduction to difference equations*. New York: Wiley, 1961.
- Groen, G., & Atkinson, R. C. Models for optimizing the learning process. *Psychological Bulletin*, 1966, 66, 309-320.
- Karush, W., & Dear, R. E. Computational studies of presentation strategies for a multilevel model of learning. *Technical Memorandum TM-3035/000/00*. Santa Monica: System Development Corporation, 1966.
- Karush, W., & Dear, R. E. Optimal stimulus presentation strategy for a stimulus sampling model of learning. *Journal of Mathematical Psychology*, 1966, 3, 19-47.
- Madden, M. S., Adams, J. A., & Spence, S. A. Memory-drum vs. adjusted-learning techniques in the study of associative interference in learning by paired-associates. *American Journal of Psychology*, 1950, 63, 186-195.
- Matheson, J. E. Optimal teaching strategies derived from mathematical learning models. Institute in Engineering-Economic Systems. *Report ccs-2*. Stanford: Stanford University, 1964.
- Osgood, C. E. *Method and theory in experimental psychology*. New York: Oxford, 1953.
- Smallwood, R. D. *A decision structure for teaching machines*. Cambridge, Mass.: MIT Press, 1962.
- Suppes, P. Problems of optimization in the learning of a simple list of items. In M. Shelly and G. Bryan (Eds.), *Human judgments and optimality*. New York: Wiley, 1964.
- Suppes, P. The uses of computers in education. *Scientific American*, 1966, 215, 206-220.
- Woodworth, R. S. A contribution to the question of "quick learning, quick forgetting." *Psychological Bulletin*, 1914, 11, 58-59.